REMARKS

This responds to the Office Action mailed on October 21, 2004.

No claims are amended, no claims are canceled, and no claims are added; as a result, claims 1-22 are now pending in this application.

§103 Rejection of the Claims

Claims 1-22 were rejected under 35 USC § 103(a) as being unpatentable over Spiegel et al. (U.S. 5,649,108) in view of Haley (U.S. 5,884,036) and Golden et al. (WO-99/53719-A1).

Applicant previously presented the following arguments for patentability.

In the Office Action the Examiner admitted that Spiegel et al does not expressly teach the requests including a request for information on the properties of the fundamental path element of a path or the properties being expressly of both the links and switches. The Examiner has held that Haley teaches receiving a request from a client for information on the properties of the fundamental path elements of a path (topology information request). However, Haley does not teach the properties being expressly of both the links and switches. The Examiner then alleges that Golden et al teaches maintaining a network map of the paths as well as the bandwidth capabilities of all links and switches along the paths (Abstract). The Examiner then concludes that it would have been obvious to a person of ordinary skill in the art at the time of the invention to include the properties of the switches in the path because this would have provided for handling changes in switch bandwidth.

However Golden teaches that the elements must be upgraded and that a reserved signaling channel is required (see page 5 of Golden). Thus combining the references would result in a method and system wherein the upgrading would be necessary. Such is not the present invention.

As describe in the present specification, the subnet manager assigns each port at least one unique address denoted a "local identification value" (LID). The subnet manager operates to discover fabric topology, assign unique addresses to all channel adapter ports that are connected

to the fabric, program switch forwarding tables, and prepare all fabric connected agents so that they can communicate with other fabric agents, in addition to performing other tasks.

According to the InfiniBand architecture specification, multiple LIDs can be assigned to each port. Each LID assigned to a port represents a unique path to this port from some other port on the cluster (or subnet). A client that wants to use multiple paths to a remote client can use different LIDs to specify different paths to its destination through the fabric. This allows a client to perform load balancing, obtain better throughput, and recover from the failure of one path if some alternate path is still functional.

To enable multi-pathing, the subnet manager identifies all possible paths to a port from any other port on the fabric, and then assigns enough LIDs to the port such that different paths to this port can be identified by a different LID. If multiple paths exist between two ports, clients on each port can explicitly choose which path will be used based on the LIDs used to communicate between the two ports. Multiple paths will exist if the subnet contains multiple (redundant) links that connect switches or channel adapters together.

Furthermore, as set forth in the present specification, the detailed path information service can be implemented as a service agent sitting on top of the General Service Interface (GSI). The General Service Interface (GSI) is an interface providing management services (e.g., connection, performance, and diagnostics) other than subnet management. Queue Pair 1 (QP1) is reserved for the GSI, which may redirect requests to other Queue Pairs (QPs).

Queries and responses to and from this service are sent using management datagrams (MADs) sent on queue pair 1. A Management Datagram (MAD) refers to the contents of an unreliable datagram packet used for communication among the HCAs, switches, routers, and TCAs to manage the network. The InfiniBand architecture specification describes the format of a number of these management commands.

The service that provides detailed path information registers with the GSI as a service agent. One option is that this service agent is implemented by the subnet administration code that also responds to the SubnAdm class of MADs. This is a natural fit since the subnet administrator is also responsible for providing other path information (like path latency, hop count, service classes supported, the maximum transfer unit and path speed) as described in the InfiniBand architecture specification. Since the MAD format for querying and reporting detailed

path information is not defined in the InfiniBand architecture specification, vendor-specific MADs can be used for this purpose.

The general header format of a vendor-specific MAD is defined in the InfiniBand architecture specification. To issue the path information query, a client would send a message with class value set to VendorSpecific; method value set to VendorSpecificGet or VendorSpecificGetTable; and attribute value set to DetailedPathInfo. This message would be sent to the subnet administrator address. If the service resides at a different local identification value (LID) or queue pair, the client can be redirected using the ClassPortInfo message specified in the InfiniBand architecture specification. As input, the client would supply relevant information like the LID or Global Identifier (GID) of the source and destination. A GID is a 128-bit identifier used to identify a port on a channel adapter, a port on a router, or a multicast group. A GID is a valid 128-bit IPv6 address (per RFC 2373) with additional properties or restrictions defined within the InfiniBand architecture specification to facilitate efficient discovery, communication, and routing.

Different implementations can also take the node GUID (Globally Unique Identifier) or platform GUID of the source and destination as input. A GUID is a software-readable number that uniquely identifies a device or component. As output of this query, the subnet administrator provides the port and node GUIDs of all switches that are traversed in this path. The width and layout of the input and output fields in the MAD are specified and documented by whoever implements the service. Note that links themselves do not have any identification or visibility and cannot be directly listed in the path information. However, the port GUIDs of the switch ports listed in the path information will uniquely identify the links being traversed. The switch port GUIDs and node GUIDs are listed in the order they are traversed from the source to the destination. For some queries, multiple packets may be needed to report the results. In this case, the mechanisms that are used to send multi-packet responses for other SubnAdm messages can be used here also.

The advantage of the foregoing arrangement is that the infrastructure in place to query and report other path properties can be used with only minor modifications to query and report detailed path information. Redirection to a different LID or queue pair can be accomplished using the infrastructure already put in place to redirect other service classes.

In the final Office Action the Examiner stated that:

"Golden was relied upon solely for its teaching a network map of all fundamental path elements including bandwidth capabilities of all links and switches along the path. When Golden performs the topology discovery (after an upgrade) is immaterial. Haley teaches the use of a request from a client for information on the properties of the fundamental path elements of a path. This is done "without modification of the path elements". The examiner further notes that Golden's topology discovery is also done "without modification of the path elements"; it just is done after there has been a modification of the path elements."

Page 5, lines 25-39 and page 6, lines 1-3 of Golden explain the required upgrading (modification) of the path elements as follows:

"As shown in Figure 2, ECP 126 communicates individually with as many switches and endstations in the LAN as have been upgraded to support the additional functionality of the present invention. For example, upgraded ones of switches 102, 114, 120 receive bandwidth reservation requests from ECP 126 and reply to ECP 126 with acknowledgments of such requests. Likewise, upgraded ones of endstations 104, 106 send connect and disconnect requests to ECP 126, which responds with acknowledgments including grants or denials of such requests. Communication between the upgraded network elements and ECP 126 is preferably performed via a reserved signaling channel 130. ECP 126 and each upgraded switch and endstation are programmed to send and recognize certain types of packets as special signaling packets."

Thus one skilled in the art would upgrade the path elements, including using a reserved signaling channel, when combining the teaching of Golden with the other cited prior art. It is clear the teachings of Golden cannot be used without modification of the path elements.

Each of the independent claims include the phrase, "<u>, the topological discovery being performed without modification of the path elements</u>". Thus, the claims distinguish over the cited prior art, taken either singly or in combination. The Examiner is therefore respectfully requested to reconsider the rejection of the claims.

Requirement for Information Under 37 C.F.R. 1.105

Pursuant to 37 C.F.R. 1.105, the Examiner has required that Applicants provide information deemed reasonably necessary for examination of the instant application.

Applicant is hereby providing a copy of the InfiniBand architecture specification (InfiniBand™ Architecture Specification, Release 1.1, November 6, 2002; http://www.infinibandta.org/specs/register/publicspec/) in an accompanying Information Disclosure Statement. See in particular pages 84-100 of volume 1.

Conclusion

Applicant respectfully submits that the claims are in condition for allowance and notification to that effect is earnestly requested. The Examiner is invited to telephone Applicant's attorney John Garrett at 847-508-2371, or the below-signed attorney at ((612) 349-9592), to facilitate prosecution of this application.

If necessary, please charge any additional fees or credit overpayment to Deposit Account No. 19-0743.

Respectfully submitted,

RAJESH R. SHAH

By his Representatives,

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
Attorneys for Intel Corporation
P.O. Box 2938
Minneapolis, Minnesota 55402
(612) 349-9592

Date April 21, 2005 By April 21, 2005
Ann M. McCrackin
Reg. No. 42,858

CERTIFICATE UNDER 37 CFR 1.8: The undersigned hereby certifies that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail, in an envelope addressed to: Mail Stop RCE, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on this 21st day of April 2005.

Chris Hammond Chris Hammond

Name Signature